Parallel Monitors for Self-adaptive Sessions*

Mario Coppo†
Università di Torino, coppo@di.unito.it

Mariangiola Dezani-Ciancaglini‡
Università di Torino, dezani@di.unito.it

Betti Venneri
Università di Firenze, venneri@unifi.it

The paper presents a data-driven model of self-adaptivity for multiparty sessions. System choreography is prescribed by a global type. Participants are incarnated by processes associated with monitors, which control their behaviour. Each participant can access and modify a set of global data, which are able to trigger adaptations in the presence of critical changes of values.
The use of the parallel composition for building global types, monitors and processes enables a significant degree of flexibility: an adaptation step can dynamically reconfigure a set of participants only, without altering the remaining participants, even if the two groups communicate.

1 Introduction

The topic of self-adaptiveness has become a key research subject within various application domains, as a response to the growing complexity of software systems operating in many different scenarios and in highly dynamic environments. To manage this complexity at a reasonable cost, novel approaches are needed in which a system can promptly react to crucial changes by reconfiguring its behaviour autonomously and dynamically, in accordance with evolving policies and objectives.

This paper proposes a *data-driven* model of *self-adaptivity* into the formal framework of *multiparty sessions* [12]. Each participant can access and modify the *control data*, whose critical values are responsible for throwing adaptation steps, in such a way that the system responds to dynamic changes by reconfiguring itself. A system comprises four active parties, as in [8]: *global types*, *monitors*, *processes*, and *adaptation functions*. Differently, the proactive role of control data in triggering adaptation is the hallmark of our approach.

A global type represents the overall communication choreography [5]; its projections onto participants generate the monitors, which set-up the communication protocols of the participants. The association of a monitor with a compliant process, dubbed *monitored process*, incarnates a participant, where the process provides the implementation of the monitoring protocol. Note that global types, monitors and processes are built using parallel composition. This is the main novelty of our calculus, which provides the system with rather flexible adaptation capabilities. Using monitors in parallel, the system is able to minimise self-adaptation by reconfiguring only some of the communications and leaving the other ones as before, in case only a part of the whole behaviour needs to be changed. Thus the adaptation of a group of participants can be transparent to another group, also in presence of communications between them, while new participants can be added and/or some of the old participants can be removed.

*Partly supported by the COST Action IC1201 BETTY.
†Partly supported by EU H2020-644235 Rephrase project, EU H2020-644298 HyVar project, ICT COST Actions IC1402 ARVI and Ateneo/CSP project RunVar.
‡Partly supported by EU H2020-644235 Rephrase project, EU H2020-644298 HyVar project, ICT COST Actions IC1402 ARVI and Ateneo/CSP project RunVar.

D. Orchard and N. Yoshida (Eds.): Programming Language Approaches to Concurrency- and Communication-Centric Software (PLACES 2016).
EPTCS 211, 2016, pp. 25–36, doi:10.4204/EPTCS.211.3

© Coppo, Dezani-Ciancaglini, & Venneri
This work is licensed under the
Creative Commons Attribution License.

The adaptation strategy is owned by control data and adaptation functions. The adaptation function contains the dynamic evolution policy, since it prescribes how the system needs to reconfigure itself, based on the changes of the data value. As long as control data do not yield critical values, the adaptation function is undefined. Thus participants communicate and read/write data according to their monitors and codes, without performing any reconfiguration.

To exemplify our approach, let us consider a company including an Italian factory (iF), an Italian supplier (iS) and a store (Ro) in Rome. Control data hold information about factories, suppliers and stores. The supplier interacts in parallel with the factory and the store about marketing data (number of item requested, delivery date). The interactions are described by the following (single threaded) global types:

$$\mathsf{G}_1 = \mu \mathbf{t}.\mathsf{iS} \to \mathsf{iF} : SF(\text{Item,Amount}).\mathsf{iF} \to \mathsf{iS} : FS(\text{DeliveryDate}).\mathbf{t}$$
$$\mathsf{G}_2 = \mu \mathbf{t}.\mathsf{Ro} \to \mathsf{iS} : RS(\text{Item,Amount}).\mathsf{iS} \to \mathsf{Ro} : SR(\text{DeliveryDate}).\mathbf{t}$$

The global type of the system is then given by $\mathbb{G}_0 = \mathsf{G}_1 \mid \mathsf{G}_2$.

The monitors of the participants are obtained as the parallel composition of the projections of these global types. For instance the monitor of participant iS is:

$$\mu \mathbf{t}.\mathsf{iF}!SF(\text{Item, Amount}).\mathsf{iF}?FS(\text{Date}).\mathbf{t} \mid \mu \mathbf{t}.\mathsf{Ro}?RS(\text{Item, Amount}).\mathsf{Ro}!SR(\text{Date}).\mathbf{t}$$

and a possible process code is [1]:

$$\mu X.y!SF(\mathsf{item},\mathsf{amount}).y?FS(date).X \mid \mu X.y?RS(item,amount).y!SR(\mathsf{date}).X$$

where ! represents output, ? represents input, and y is a channel.

Assume that the control data reveals the opening of a new store (Lo) in London, which must receive its items from the Italian seller.

A new global type is obtained by putting in parallel to the current global type the type G, produced by applying the adaptation function to the control data:

$$\mathsf{G} = \mu \mathbf{t}.\mathsf{Lo} \to \mathsf{iS} : LS(\text{Item,Amount}).\mathsf{iS} \to \mathsf{Lo} : SL(\text{DeliveryDate}).\mathbf{t}$$

The global type G adds the participant Lo to the conversation and modifies the participant iS by adding the monitor $\mu \mathbf{t}.\mathsf{Lo}?LS(\text{Item, Amount}).\mathsf{Lo}!SL(\text{Date}).\mathbf{t}$ and the process

$$\mu X.y?LS(item,amount).y!SL(\mathsf{date}).X$$

in parallel with his previous ones.

Outline This paper has a standard structure. After a further example (Section 2) we present syntax (Section 3), types (Section 4) and semantics (Section 5) of our calculus. In Section 6 we draw some conclusions and discuss related works.

2 An Extended Example

Let us now extend the example of the Introduction by considering a company Ada in which factories interact with the general manager about production policies and suppliers interact with factories and stores by exchanging marketing data. Moreover, the stores can communicate each other for requiring some products. The company initially consists of:

- a general manager (GM);
- an Italian (iF) and an American Factory (aF);
- an Italian (iS) and an American Supplier (aS);
- three stores, located in Rome (Ro), New York (NY) and Chicago (Ch).

The global type prescribing the communications is $\mathbb{G} = \mathsf{G}_1 \mid \mathsf{G}_2 \mid \mathsf{G}_3 \mid \mathsf{G}_4 \mid \mathsf{G}_5 \mid \mathsf{G}_6$, where $\mathsf{G}_1, \mathsf{G}_2$ are as in the Introduction and:

[1] Sorts are written with upper case initials, expression *variables* and values with lower case initials, but different fonts.

$$\mathsf{G}_3 = \mu\mathbf{t}.\mathsf{GM} \to \mathsf{iF} : GIF(\text{ProductionLines}).\mathsf{GM} \to \mathsf{aF} : GAF(\text{ProductionLines}).$$
$$\mathsf{iF} \to \mathsf{GM} : IFG(\text{ProgressReport}).\mathsf{aF} \to \mathsf{GM} : AFG(\text{ProgressReport}).\mathbf{t}$$
$$\mathsf{G}_4 = \mu\mathbf{t}.\mathsf{aS} \to \mathsf{aF} : SF(\text{Item,Amount}).\mathsf{aF} \to \mathsf{aS} : FS(\text{DeliveryDate}).\mathbf{t}$$
$$\mathsf{G}_5 = \mu\mathbf{t}.\mathsf{Ch} \to \mathsf{aS} : CS(\text{Item,Amount}).\mathsf{aS} \to \mathsf{Ch} : SC(\text{DeliveryDate}).\mathbf{t}$$
$$\mathsf{G}_6 = \mu\mathbf{t}.\mathsf{NY} \to \mathsf{aS} : NS(\text{Item,Amount}).$$
$$\mathsf{aS} \to \mathsf{NY} : \{YES(\text{DeliveryDate}).\mathsf{NY} \to \mathsf{Ch} : NCY(\text{NoItem}).\mathbf{t},$$
$$NO(\text{NoItem}) : \mathsf{NY} \to \mathsf{Ch} : NCN(\text{Item,Amount}).\mathsf{Ch} \to \mathsf{NY} : CN(\text{DeliveryDate}).\mathbf{t}\}$$

Note that the New York store can ask for items both to the America factory and to the Chicago store. Notice also that G_1 and G_4 differ for the participants, but not for the labels, and this allows the process

$$\mu X.y?SF(item, amount).y!FS(\mathsf{deliveryDate}).X$$

to incarnate both iF for G_1 and aF for G_4. This process in parallel with

$$\mu X.y?GIF(\text{ProductionLines}).y!IFG(\text{ProgressReport}).X,$$

where y must be replaced by the appropriate run time channel, plays the role of iF for the whole $\mathbb{G}$.

Assume now that the catastrophic event of a fire, incapacitating the American factory, requires the Ada company to update itself. The American factory modifies the control data by putting the information about its unavailability. The function F, applied to the modified control data, kills aF and produces the global type:

$$\mathbb{G}' = \mathsf{FF} \to \mathsf{GM} : FG(\text{DamagesReport}).$$
$$\mu\mathbf{t}.\mathsf{GM} \to \mathsf{TS} : GT(\text{ReconstructionPlan}).\mathsf{TS} \to \mathsf{GM} : TG(\text{TechnicalRevisions}).\mathbf{t} \mid$$
$$\mu\mathbf{t}.\mathsf{aS} \to \mathsf{iF} : ASF(\text{Item,Amount}).\mathsf{iF} \to \mathsf{aS} : FAS(\text{DeliveryDate}).\mathbf{t} \mid$$
$$\mu\mathbf{t}.\mathsf{GM} \to \mathsf{iF} : GIF(\text{ProductionLines}).\mathsf{iF} \to \mathsf{GM} : IFG(\text{ProgressReport}).\mathbf{t}$$

According to $\mathbb{G}'$, the FireFighters (FF) send to the general manager a report on the damages and then the general manager opens with the Technical Service (TS) a conversation about the reconstruction plan. The American seller now asks for products to the Italian factory. Notice that the labels of the communications between them must be different from those in G_1, which is not modified, since the Italian factory continues to serve the Italian seller as before. The new global type is the parallel of $\mathbb{G}'$ with the global type obtained from $\mathbb{G}$ by erasing G_3 (in which all communications involve aF or they are between GM and iF) and erasing G_4 (in which all communications involve aF). The rational is to erase the communications in which

- either one of the two participants is killed,
- or both participants are reconfigured.

The new process incarnating iF is obtained by putting

$$\mu X.y?ASF(\text{item,amount}).y!FAS(\text{deliveryDate}).X \mid \mu X.y?GIF(\text{ProductionLines}).y!IFG(\text{ProgressReport}).X$$

in parallel with the previous process incarnating iF. Then the code of the iF participant becomes:

$$\mu X.y?SF(item, amount).y!FS(\mathsf{deliveryDate}).X \mid \mu X.y?ASF(item, amount).y!FAS(\mathsf{deliveryDate}).X \mid$$
$$\mu X.y?GIF(\text{ProductionLines}).y!IFG(\text{ProgressReport}).X$$

where y must be replaced by the appropriate run time channel.

For readability in this example (and in that of the Introduction) we described adaptation as modification of global types. In the semantics (Section 5) global types are left implicit and the adaptation rule modifies monitors and processes.

3 Syntax

3.1 Global Types

Following a widely common approach [5], the set-up of protocols starts from global types. Global types establish overall communication schemes. In our setting they also control the reconfiguration phase, in

$$(\mathsf{p} \to \mathsf{q} : \{\ell_i(S_i).\mathsf{G}_i\}_{i\in I}) \upharpoonright \mathsf{r} = \begin{cases} \mathsf{p}?\{\ell_i(S_i).\mathsf{G}_i \upharpoonright \mathsf{q}\}_{i\in I} & \text{if } \mathsf{r} = \mathsf{q} \\ \mathsf{q}!\{\ell_i(S_i).\mathsf{G}_i \upharpoonright \mathsf{r}\}_{i\in I} & \text{if } \mathsf{r} = \mathsf{p} \\ \mathsf{G}_{i_0} \upharpoonright \mathsf{r} & \text{where } i_0 \in I \text{ if } \mathsf{r} \neq \mathsf{p} \text{ and } \mathsf{r} \neq \mathsf{q} \\ & \text{and } \mathsf{G}_i \upharpoonright \mathsf{r} = \mathsf{G}_j \upharpoonright \mathsf{r} \text{ for all } i, j \in I \\ \mathsf{u}?\bigcup_{k\in K}\{\ell_k(S_k).\mathsf{G}_k \upharpoonright \mathsf{r}\} & \text{where } K = \bigcup_{i\in I} J_i \text{ if } \mathsf{r} \neq \mathsf{p},\ \mathsf{r} \neq \mathsf{q} \text{ and} \\ & \mathsf{G}_i \upharpoonright \mathsf{r} = \mathsf{u}?\{\ell_j(S_j).\mathsf{G}_j\}_{j\in J_i} \text{ for all } i \in I \text{ and} \\ & J_m \cap J_n = \emptyset \text{ and } \ell_m \neq \ell_n \\ & \text{for all } m, n \in K \text{ such that } m \neq n \end{cases}$$

$$(\mu \mathbf{t}.\mathsf{G}) \upharpoonright \mathsf{p} = \begin{cases} \mu \mathbf{t}.\mathsf{G} \upharpoonright \mathsf{p} & \text{if } \mathsf{p} \in \mathsf{G}, \\ \mathsf{end} & \text{otherwise.} \end{cases} \qquad \mathbf{t} \upharpoonright \mathsf{p} = \mathbf{t} \qquad \mathsf{end} \upharpoonright \mathsf{p} = \mathsf{end} \qquad (\mathbb{G} \mid \mathbb{G}') \upharpoonright \mathsf{p} = (\mathbb{G} \upharpoonright \mathsf{p}) \mid (\mathbb{G}' \upharpoonright \mathsf{p})$$

Table 1: Projection of a global type onto a participant.

which a system adapts itself to new environmental conditions.

Let L be a set of *labels*, ranged over by ℓ, which mark the exchanged values as in [10]. We assume some basic *sorts*, ranged over by S, i.e. $S ::= \mathsf{Bool} \mid \mathsf{Int} \mid \ldots$.

Single threaded global types give sequential communications with alternatives.
In $\mathsf{p} \to \mathsf{q} : \{\ell_i(S_i).\mathsf{G}_i\}_{i\in I}$ participant p sends to participant q a label l_i together with a value of sort S_i for some $i \in I$. We implicitly assume that $\ell_i \neq \ell_j$ for all $i \neq j$.
We take an equi-recursive view of global types, monitors, processes and process types, identifying $\mu \mathbf{t}.\mathsf{G}$ with $\mathsf{G}\{\mu \mathbf{t}.\mathsf{G}/\mathbf{t}\}$ etc. We assume that all recursions are guarded. In writing examples we omit brackets when there is only one branch.

Global types are parallel compositions of single threaded global types with distinct labels for common participants.

Definition 3.1 *1.* Single threaded global types *are defined by:*

$$\mathsf{G} \ ::= \ \mathsf{p} \to \mathsf{q} : \{\ell_i(S_i).\mathsf{G}_i\}_{i\in I} \mid \mu t.\mathsf{G} \mid t \mid \mathsf{end}$$

2. Global types *are defined by:* $\mathbb{G} \ ::= \ \mathsf{G} \mid \mathbb{G} \mid \mathbb{G}$, *where global types in parallel have distinct labels for common participants.*

We allow parallels of global types with shared participants, a possibility usually forbidden [12, 5, 3].

3.2 Monitors

Single threaded monitors can be viewed as projections of single threaded global types onto individual participants, as in the standard approach of [12] and [2]. In our calculus, however, monitors are more than types, as in [8]: they have an active role in system dynamics, since they guide communications and adaptations. As expected monitors are parallel compositions of single threaded monitors with distinct labels.

Definition 3.2 *1.* Single threaded monitors *are defined by:*

$$\mathsf{M} \ ::= \ \mathsf{p}?\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I} \mid \mathsf{q}!\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I} \mid \mu t.\mathsf{M} \mid t \mid \mathsf{end}$$

2. Monitors *are defined by:* $\mathbb{M} \ ::= \ \mathsf{M} \mid \mathbb{M} \mid \mathbb{M}$, *where monitors in parallel have distinct labels.*

An input monitor $\mathsf{p}?\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I}$ is able to drive a process that can receive, for each $i \in I$, a value of sort S_i, labeled by ℓ_i, having as continuation a process which is adequate for M_i. This corresponds to an

external choice. Dually an output monitor $\mathsf{q}!\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I}$ is able to drive a process which can send (by an internal choice) a value of sort S_i, labeled by ℓ_i, and then continues as prescribed by M_i for each $i \in I$.

The projection of global types onto participants is given in Table 1. A projection onto a participant r not involved, as sender or receiver, in a choice is defined if either the projection onto r of all continuations are the same (condition $\mathsf{G}_i \upharpoonright \mathsf{q} = \mathsf{G}_j \upharpoonright \mathsf{q}$ for all $i, j \in I$) or all these projections can be merged in an input monitor (last case), as in [10]. A global type G is *well formed* if its projections are defined for all participants. We assume always that global types are well formed.

3.3 Processes

Processes represent code that is associated to monitors in order to implement participants. As in [8] and differently from standard session calculi (see [13] and the references there), processes do not specify the participants involved in sending and receiving actions. The associated monitors determine senders and receivers.

The communication actions of processes are performed through *channels*. Each process owns a unique channel endpoint. We use y to denote this channel endpoint in the process code. As usual, the user channel y will be replaced at run time by a session channel $\mathsf{s}[\mathsf{p}]$ (where s is the session name and p is the current participant). Let c denote a user channel or a session channel.

Definition 3.3 *1.* Single threaded processes *are defined by:*

$$P \quad ::= \quad \mathsf{c}?\ell(x).P \;\mid\; \mathsf{c}!\ell(\mathsf{e}).P \;\mid\; P+P \;\mid\; \mathsf{if\ e\ then}\ P\ \mathsf{else}\ P \;\mid\; \mathsf{op}.P \;\mid\; \mu X.P \;\mid\; X \;\mid\; \mathbf{0}$$

2. Processes *are defined by:* $\mathbb{P} \quad ::= \quad P \quad \mid \mathbb{P} \mid \mathbb{P}$.

The syntax of processes is rather standard, in particular the operator $+$ represents external choice. A system has control data (see Definition 3.5) and the op operator represents an action on this data, for instance a "read" or "write" operation. We leave unspecified the kind of actions, since we are only interested in the dynamic changes of this data, which determine the self-reconfiguration of the whole system.

3.4 Networks

A process is always controlled by a monitor, which ensures that all performed actions fit the protocol prescribed by the global type. Each monitor controls a single process. So participants correspond to pairs of processes and monitors. We write $\mathbb{M}[\mathbb{P}]$ to represent a process $\mathbb{P}$ controlled by a monitor $\mathbb{M}$, dubbed *monitored process*. The sessions are initiated by the "new" constructor applied to a global type (*session initiator*), denoted by $\mathsf{new}(\mathbb{G})$, which generates the monitors and associates them with adequate processes (see Definition 4.2). The parallel composition of session initiators and processes with the corresponding monitors form a network. Networks can be restricted on session names.

Definition 3.4 Networks *are defined by:* $N \;::=\; \mathsf{new}(\mathbb{G}) \;\mid\; \mathbb{M}[\mathbb{P}] \;\mid N \parallel N \;\mid\; (\nu \mathsf{s})N$.

By $\mathbb{M}_\mathsf{p}[\mathbb{P}_\mathsf{p}]$ we denote the monitored process of participant p, i.e. the channel in $\mathbb{P}_\mathsf{p}$ must be $\mathsf{s}[\mathsf{p}]$ for some s. We use Π to represent the parallel composition of monitored process.

3.5 Systems

A system includes a network, a global state and assumes a collection of processes denoted by $\mathscr{P}$. A main component of the global state is the *adaptation function* F, which says how to run adaptations according to critical data variations. The global state contains also *control data* σ, which are updated by processes. We represent systems as the composition (via "$\between$") of a network and control data. We omit the adaptation function and the process collection, considering them as implicit parameters, since they are not modified by reductions.

$$\Gamma \vdash \mathbf{0} \rhd \mathsf{c} : \mathsf{end}\ \text{END} \qquad \frac{\Gamma \vdash P \rhd \mathsf{c} : \mathsf{T}}{\Gamma \vdash \mathsf{op}.P \rhd \mathsf{c} : \mathsf{T}}\ \text{OP} \qquad \Gamma, X : \mathsf{T} \vdash X \rhd \mathsf{c} : \mathsf{T}\ \text{AX}$$

$$\frac{\Gamma, X : \mathsf{T} \vdash P \rhd \mathsf{c} : \mathsf{T}}{\Gamma \vdash \mu X.P \rhd \mathsf{c} : \mathsf{T}}\ \text{REC} \qquad \frac{\Gamma, x : S \vdash P \rhd \mathsf{c} : \mathsf{T}}{\Gamma \vdash \mathsf{c}?\ell(x).P \rhd \mathsf{c} :?\ell(S).\mathsf{T}}\ \text{RCV} \qquad \frac{\Gamma \vdash P \rhd \mathsf{c} : \mathsf{T} \qquad \Gamma \vdash \mathsf{e} : S}{\Gamma \vdash \mathsf{c}!\ell(\mathsf{e}).P \rhd \mathsf{c} :!\ell(S).\mathsf{T}}\ \text{SEND}$$

$$\frac{\Gamma \vdash P_1 \rhd \mathsf{c} : \mathsf{T}_1 \quad \Gamma \vdash P_2 \rhd \mathsf{c} : \mathsf{T}_2 \quad \mathsf{T}_1 \wedge \mathsf{T}_2 \in \mathscr{T}}{\Gamma \vdash P_1 + P_2 \rhd \mathsf{c} : \mathsf{T}_1 \wedge \mathsf{T}_2}\ \text{CHOICE}$$

$$\frac{\Gamma \vdash \mathsf{e} : \mathsf{Bool} \quad \Gamma \vdash P_1 \rhd \mathsf{c} : \mathsf{T}_1 \quad \Gamma \vdash P_2 \rhd \mathsf{c} : \mathsf{T}_2 \quad \mathsf{T}_1 \vee \mathsf{T}_2 \in \mathscr{T}}{\Gamma \vdash \mathsf{if}\ \mathsf{e}\ \mathsf{then}\ P_1\ \mathsf{else}\ P_2 \rhd \mathsf{c} : \mathsf{T}_1 \vee \mathsf{T}_2}\ \text{IF}$$

$$\frac{\Gamma \vdash \mathbb{P}_1 \rhd \mathsf{c} : \mathbb{T}_1 \quad \Gamma \vdash \mathbb{P}_2 \rhd \mathsf{c} : \mathbb{T}_2}{\Gamma \vdash \mathbb{P}_1 \mid \mathbb{P}_2 \rhd \mathsf{c} : \mathbb{T}_1 \mid \mathbb{T}_2}\ \text{PAR}$$

Table 2: Typing rules for processes.

Definition 3.5 Systems *are defined by:* $\mathscr{S} ::= N \between \sigma$.

4 Type System

Process types describe process communication behaviours [11]. They have prefixes corresponding to sending and receiving of labels and values.

Definition 4.1 *1.* Single threaded types *are inductively defined by:*

$$\mathsf{T} ::= \bigwedge_{i\in I} ?\ell_i(S_i).\mathsf{T}_i \mid \bigvee_{i\in I} !\ell_i(S_i).\mathsf{T}_i \mid \mu \boldsymbol{t}.\mathsf{T} \mid \boldsymbol{t} \mid \mathsf{end}$$

where all labels in intersections and unions are different.

2. Types *are inductively defined by:* $\mathbb{T} ::= \mathsf{T} \mid \mathbb{T} \mid \mathbb{T}$, *where types in parallel have distinct labels.*

Types are built from input and output prefixes ($?l(S)$ and $!l(S)$) by means of constructs for intersection types (used to type external choices) and union types (used to type conditionals).

An *environment* Γ is a finite mapping from expression variables to sorts and from process variables to single threaded types: $\Gamma ::= \emptyset \mid \Gamma, x : S \mid \Gamma, X : \mathsf{T}$, where the notation $\Gamma, x : S$ ($\Gamma, X : \mathsf{T}$) means that x (X) does not occur in Γ.

We assume that expressions are typed by sorts, as usual. The typing judgments for expressions are of the shape $\Gamma \vdash \mathsf{e} : S$ and the typing rules for expressions are standard.

Single threaded processes can be typed only if they contain at most one channel. This choice is justified by the design of monitored processes as session participants and by the absence of delegation. Therefore the typing judgments for single threaded processes have the form $\Gamma \vdash P \rhd \mathsf{c} : \mathsf{T}$ and for processes $\Gamma \vdash \mathbb{P} \rhd \mathsf{c} : \mathbb{T}$. Typing rules for processes are given in Table 2. The external choice is typed by an intersection type, since an external choice offers both behaviours of the composing processes. Dually, a conditional is an internal choice and so it is typed by a union type.

As in [8], a single threaded process is equipped with a single threaded type, ranged over by T, describing its communication actions. A parallel composition of single threaded processes is typed by the parallel composition of the corresponding single threaded types. We use $\mathbb{T}$ to range over types and we write $\vdash \mathbb{P} \rhd \mathsf{c} : \mathbb{T}$ if process $\mathbb{P}$ has the unique channel c typed by $\mathbb{T}$.

$$[\text{SUB-END}]\quad \mathsf{end} \leq \mathsf{end} \qquad [\text{SUB-IN}]\quad \frac{\forall i \in I: \quad \mathsf{T}_i \leq \mathsf{T}'_i}{\bigwedge_{i\in I\cup J} \mathsf{p}?\ell_i(S_i).\mathsf{T}_i \leq \bigwedge_{i\in I} \mathsf{p}?\ell_i(S_i).\mathsf{T}'_i} \qquad [\text{SUB-OUT}]\quad \frac{\forall i \in I: \quad \mathsf{T}_i \leq \mathsf{T}'_i}{\bigvee_{i\in I} \mathsf{p}!\ell_i(S_i).\mathsf{T}_i \leq \bigvee_{i\in I\cup J} \mathsf{p}!\ell_i(S_i).\mathsf{T}'_i}$$

Table 3: Subtyping.

We end this section defining *adequacy* between a process $\mathbb{P}$ and a monitor $\mathbb{M}$, denoted $\mathbb{P} \propto \mathbb{M}$, which is based on the matching between types and monitors. This matching is made rather flexible by using the *subtype* relation on types defined in Table 3. The double line in rules indicates that the rules are interpreted *coinductively* [14, 21.1]. Subtyping is monotone, for input/output prefixes, with respect to continuations and it follows the usual set theoretic inclusion of intersection and union.

Definition 4.2 *1. A single threaded process P is* adequate *for a single threaded monitor* M *(notation $P \propto \mathsf{M}$) if $\vdash P \rhd \mathsf{c} : \mathsf{T}$ and $\mathsf{T} \leq |\mathsf{M}|$, where the mapping $|\ |$ is defined by:*

$$|\mathsf{p}?\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I}| = \bigwedge_{i\in I} \mathsf{p}?\ell_i(S_i).|\mathsf{M}_i| \qquad |\mathsf{q}!\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I}| = \bigvee_{i\in I} \mathsf{q}!\ell_i(S_i).|\mathsf{M}_i|$$

$$|\mu \boldsymbol{t}.\mathsf{M}| = \mu \boldsymbol{t}.|\mathsf{M}| \qquad |\boldsymbol{t}| = \boldsymbol{t} \qquad |\mathsf{end}| = \mathsf{end}$$

2. The adequacy *of a process $\mathbb{P}$ for a monitor $\mathbb{M}$ (notation $\mathbb{P} \propto \mathbb{M}$) is the smallest relation such that:*

- $P \propto \mathsf{M}$*;*
- $\mathbb{P}_1 \propto \mathbb{M}_1$ *and* $\mathbb{P}_2 \propto \mathbb{M}_2$ *imply* $\mathbb{P}_1 \mid \mathbb{P}_2 \propto \mathbb{M}_1 \mid \mathbb{M}_2$.

Note that adequacy is decidable, since processes have unique types and subtyping is decidable.

The following lemma states a crucial property of the adequacy relation. A process $\mathbb{P}$ is adequate for a monitor $\mathbb{M}$ iff $\mathbb{P}$ is the parallel composition of single threaded processes, $\mathbb{M}$ is the parallel composition of single threaded monitors and each single threaded process is adequate for exactly one single threaded monitor and vice versa.

Lemma 4.3 *If $\mathbb{P} \propto \mathbb{M}$, then $\mathbb{P} = \Pi_{i\in I} P$ and $\mathbb{M} = \Pi_{i\in I} \mathsf{M}$ such that $P_i \propto \mathsf{M}_i$ for all $i \in I$ and $P_i \not\propto \mathsf{M}_j$ for all $i, j \in I$ with $i \neq j$.*

Under the hypothesis of previous lemma we use $\mathbb{P}.i$ and $\mathbb{M}.i$ to denote P_i and M_i, respectively.

5 Semantics

Processes can communicate labels and values, or can read/modify the control data trough op operations. The semantics of processes is described via a LTS defined in Table 4, where the treatment of recursions and conditionals is standard.

In the rules for external choice, β ranges over $\mathsf{s}[\mathsf{p}]?\ell(v), \mathsf{s}[\mathsf{p}]!\ell(v)$ and γ ranges over op, τ. In the rule for parallel composition δ ranges over β and γ. We omit the symmetric rules. The external choices

$$\mathsf{op}.P \xrightarrow{\mathsf{op}} P \qquad \mu X.P \xrightarrow{\tau} P\{\mu X.P/X\} \qquad \mathsf{s}[\mathsf{p}]?\ell(x).P \xrightarrow{\mathsf{s}[\mathsf{p}]?\ell(v)} P\{v/x\} \qquad \mathsf{s}[\mathsf{p}]!\ell(\mathsf{e}).P \xrightarrow{\mathsf{s}[\mathsf{p}]!\ell(v)} P \quad \mathsf{e}\downarrow v$$

$$\text{if } \mathsf{e} \text{ then } P \text{ else } Q \xrightarrow{\tau} P \quad \mathsf{e}\downarrow \mathsf{true} \qquad\qquad \text{if } \mathsf{e} \text{ then } P \text{ else } Q \xrightarrow{\tau} Q \quad \mathsf{e}\downarrow \mathsf{false}$$

$$\frac{P \xrightarrow{\beta} P'}{P + Q \xrightarrow{\beta} P'} \qquad \frac{P \xrightarrow{\gamma} P'}{P + Q \xrightarrow{\gamma} P' + Q} \qquad \frac{\mathbb{P} \xrightarrow{\delta} \mathbb{P}'}{\mathbb{P} \mid \mathbb{P}'' \xrightarrow{\delta} \mathbb{P}' \mid \mathbb{P}''}$$

Table 4: LTS of processes.

are done by the communication actions, while the operations on the global state are transparent. This is needed since the operations on the control data are recorded neither in the process types nor in the monitors. An operation on data in an external choice can be performed also if a branch, different from that containing the operation, is then executed.

We assume a standard structural equivalence on processes, monitors, and networks, in which the parallel operator is commutative and associative, and $\mathbf{0}$ and end are neutral elements. Moreover, the process $\mathbf{0}$ with end monitor behaves as the neutral element for the parallel of networks and it absorbs restrictions, that is: $\mathsf{end}[\mathbf{0}] \mid N \equiv N$ and $(\nu \mathsf{s})\mathsf{end}[\mathbf{0}] \equiv \mathsf{end}[\mathbf{0}]$.

The evolution of a system depends on the evolution of its network and control data. The basic components of networks are the openings of sessions (through the new operator on global types) and the processes associated with monitors.

A session starts by reducing a network $\mathsf{new}(\mathbb{G})$. If $\mathbb{G}$ is a single threaded global type, then for each p in the set $\mathsf{pa}(\mathsf{G})$ of the participants in the global type G, we need to find a process P_p in the collection $\mathscr{P}$ associated to the current system. The process P_p must be adequate for the monitor which is the projection of G onto p. Then the process (where the channel y has been replaced by $s[\mathsf{p}]$) is associated to the corresponding monitor. Lastly, the name s is restricted.

$$\frac{\mathsf{M}_\mathsf{p} = \mathsf{G} \upharpoonright \mathsf{p} \quad \forall \mathsf{p} \in \mathsf{pa}(\mathsf{G}).\ P_\mathsf{p} \in \mathscr{P} \ \&\ P_\mathsf{p} \propto \mathsf{M}_\mathsf{p}}{\mathsf{new}(\mathsf{G}) \longrightarrow (\nu \mathsf{s})\ (\prod_{\mathsf{p} \in \mathsf{pa}(\mathsf{G})} \mathsf{M}_\mathsf{p}[P_\mathsf{p}\{\mathsf{s}[\mathsf{p}]/y\}])}\ \textsc{InitS}$$

The new applied to the parallel of two global types simply puts in parallel the networks produced by the news of the two global types, by restricting them with the same session name. In this way we ensure the privacy of the communications in a session (as standard in session calculi [12]).

$$\frac{\mathsf{new}(\mathbb{G}) \longrightarrow (\nu \mathsf{s})\ N \quad \mathsf{new}(\mathbb{G}') \longrightarrow (\nu \mathsf{s})\ N'}{\mathsf{new}(\mathbb{G} \mid \mathbb{G}') \longrightarrow (\nu \mathsf{s})\ (N \mid N')}\ \textsc{InitM}$$

Monitors guide the communications of processes by choosing the senders/receivers and by allowing only some actions among those offered by the processes. This is formalised by the following LTS for monitors:

$$\mathsf{p}?\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I} \xrightarrow{\mathsf{p}?\ell_j} \mathsf{M}_j \quad j \in I \qquad \mathsf{q}!\{\ell_i(S_i).\mathsf{M}_i\}_{i\in I} \xrightarrow{\mathsf{q}!\ell_j} \mathsf{M}_j \quad j \in I \qquad \frac{\mathsf{M} \xrightarrow{\alpha} \mathsf{M}'}{\mathsf{M} \mid \mathbb{M}'' \xrightarrow{\alpha} \mathsf{M}' \mid \mathbb{M}''}$$

where α ranges over $\mathsf{p}?l$ and $\mathsf{q}!l$.

Rule [Com] allows monitored processes to exchange messages:

$$\frac{\mathbb{M}_1 \xrightarrow{\mathsf{q}?\ell} \mathbb{M}_1' \quad \mathbb{P}_1 \xrightarrow{\mathsf{s}[\mathsf{p}]?\ell(v)} \mathbb{P}_1' \quad \mathbb{M}_2 \xrightarrow{\mathsf{p}!\ell} \mathbb{M}_2' \quad \mathbb{P}_2 \xrightarrow{\mathsf{s}[\mathsf{q}]!\ell(v)} \mathbb{P}_2'}{\mathbb{M}_1[\mathbb{P}_1] \parallel \mathbb{M}_2[\mathbb{P}_2] \longrightarrow \mathbb{M}_1'[\mathbb{P}_1'] \parallel \mathbb{M}_2'[\mathbb{P}_2']}\ \textsc{Com}$$

Monitored processes can modify control data thanks to rule [OP]:

$$\frac{\mathbb{P} \xrightarrow{\mathsf{op}} \mathbb{P}'}{\mathbb{M}[\mathbb{P}] \between \sigma \longrightarrow \mathbb{M}[\mathbb{P}'] \between \mathsf{op}(\sigma)}\ \textsc{OP}$$

Adaptation is triggered by a function which applied to the state σ returns a global type $\mathbb{G}$ which represents the choreography of the new part of the system, a set $\mathscr{K}$ of participants that are removed and a new state σ'. The outcome of the adaptation function determines a reconfiguration of the system as summarised below. Let $\mathscr{A}$ denote the set of the participants before adaptation. The system running before the adaptation step is modified in the following way:

1. the participants in $\mathscr{K}$ are removed;
2. all communications between participants in $\mathscr{A} \setminus \mathscr{K}$ and participants in $\mathscr{K}$ are erased from all monitors of participants in $\mathscr{A} \setminus \mathscr{K}$;
3. all communications between two participants in $\mathscr{A} \cap \mathsf{pa}(\mathbb{G})$ are erased from all monitors of participants in $\mathscr{A} \cap \mathsf{pa}(\mathbb{G})$.
4. The monitors resulting as projections of $\mathbb{G}$ are added, in parallel with the monitors of participants in $\mathscr{A} \cap \mathsf{pa}(\mathbb{G})$ obtained as described in (3), or as monitors of new participants in $\mathsf{pa}(\mathbb{G}) \setminus \mathscr{A}$.
5. The processes with adapted monitors are modified in order to be adequate to these monitors.

To define the formal rule for adaptation we need to define the mappings mon and proc. By $\mathbb{M} \setminus \mathscr{A}$ we denote the monitor obtained from $\mathbb{M}$ by erasing all communications to participants belonging to $\mathscr{A}$.

$$\mathsf{mon}(\mathsf{p}, \mathbb{M}, \mathbb{G}, \mathscr{K}) = \begin{cases} \mathbb{M} \setminus \mathscr{K} & \text{if } \mathsf{p} \notin \mathsf{pa}(\mathbb{G}), \\ (\mathbb{G} \upharpoonright \mathsf{p}) \mid (\mathbb{M} \setminus (\mathsf{pa}(\mathbb{G}) \cup \mathscr{K})) & \text{otherwise.} \end{cases}$$

The mapping mon applied to a participant p, his current monitor $\mathbb{M}$, a global type $\mathbb{G}$ and a set of killed participants $\mathscr{K}$ gives the new monitor for p. If p is not a participant of $\mathbb{G}$, then the new monitor is simply $\mathbb{M}$ where all communications with killed participants are erased. Note that adaptation is transparent to p whenever p does not communicate with killed participants. Otherwise the new monitor is the parallel composition of the projection of $\mathbb{G}$ onto p with the monitor $\mathbb{M}$, where all communications with participants of $\mathbb{G}$ and killed participants are erased. Clearly mon is a partial mapping, since it is undefined when some labels occur both in $\mathbb{G} \upharpoonright \mathsf{p}$ and in $\mathbb{M} \setminus (\mathsf{pa}(\mathbb{G}) \cup \mathscr{K})$. We could make mon total simply by always performing a fresh renaming of the labels in $\mathbb{G}$. This dramatic solution however would never allow us to use the current process for a participant in $\mathsf{pa}(\mathbb{G})$, even if the process is adequate for the new monitor. For this reason in rule [ADAPT] by $\hat{\mathbb{G}}$ we denote $\mathbb{G}$ if $\mathsf{mon}(\mathsf{p}, \mathbb{M}, \mathbb{G}, \mathscr{K})$ is defined for all $\mathsf{pa}(\mathbb{G})$ and a fresh relabelling of $\mathbb{G}$ otherwise.

$$\mathsf{procS}(\mathsf{s}[\mathsf{p}], P, \mathsf{M}, \mathscr{A}) = \begin{cases} P & \text{if } P \propto \mathsf{M} \setminus \mathscr{A}, \\ Q\{\mathsf{s}[\mathsf{p}]/y\} & \text{if } P \not\propto \mathsf{M} \setminus \mathscr{A} \ \&\ Q \in \mathscr{P} \ \&\ Q \propto \mathsf{M} \setminus \mathscr{A}. \end{cases}$$

The mapping procS deals with a single threaded process P (with channel $\mathsf{s}[\mathsf{p}]$), a single threaded monitor M and a set $\mathscr{A}$ of processes. If P is adequate for M without the communications to participants in $\mathscr{A}$, i.e. for $\mathsf{M} \setminus \mathscr{A}$, then the mapping returns P. Otherwise the mapping takes a process Q (adequate for $\mathsf{M} \setminus \mathscr{A}$) in the collection $\mathscr{P}$ and it returns Q where y is replaced by $\mathsf{s}[\mathsf{p}]$, i.e. $Q\{\mathsf{s}[\mathsf{p}]/y\}$.

$$\mathsf{proc}(\mathsf{s}[\mathsf{p}], \mathbb{P}, \mathbb{M}, \mathbb{G}, \mathscr{K}) = \begin{cases} \mathbb{P} & \text{if } \mathbb{P} \propto \mathsf{mon}(\mathsf{p}, \mathbb{M}, \mathbb{G}, \mathscr{K}) \\ \Pi_{i \in I} \mathsf{procS}(\mathsf{s}[\mathsf{p}], \mathbb{P}.i, \mathbb{M}.i, \mathscr{K}) & \text{if } \mathsf{p} \notin \mathsf{pa}(\mathbb{G}) \\ \mathbb{Q}\{\mathsf{s}[\mathsf{p}]/y\} \mid \mathbb{P}' & \text{if } \mathsf{p} \in \mathsf{pa}(\mathbb{G}) \ \&\ \mathbb{Q} \in \mathscr{P} \ \&\ \mathbb{Q} \propto \mathbb{G} \upharpoonright \mathsf{p} \end{cases}$$
$$\text{where } \mathbb{P}' = \Pi_{i \in I} \mathsf{procS}(\mathsf{s}[\mathsf{p}], \mathbb{P}.i, \mathbb{M}.i, \mathsf{pa}(\mathbb{G}) \cup \mathscr{K}).$$

The mapping proc returns $\mathbb{P}$ whenever $\mathbb{P}$ is adequate for $\mathsf{mon}(\mathsf{p}, \mathbb{M}, \mathbb{G}, \mathscr{K})$. Otherwise both $\mathbb{P}$ and $\mathbb{M}$ are split in the corresponding single threaded processes ($\mathbb{P}.i$) and monitors ($\mathbb{M}.i$) for $i \in I$ according to Lemma 4.3. Each $\mathbb{P}.i$ and $\mathbb{M}.i$ for $i \in I$ become arguments of procS. If $\mathsf{p} \notin \mathsf{pa}(\mathbb{G})$ the set of processes argument of procS is $\mathscr{K}$. Otherwise the set of processes argument of procS is $\mathsf{pa}(\mathbb{G}) \cup \mathscr{K}$. This reflects the two cases in the definition of mon. Moreover if $\mathsf{p} \in \mathsf{pa}(\mathbb{G})$ the final parallel contains also a process $\mathbb{Q}$ in $\mathscr{P}$ (adequate for $\mathbb{G} \upharpoonright \mathsf{p}$) with y replaced by $\mathsf{s}[\mathsf{p}]$.
For example $\mathsf{proc}(\mathsf{s}[\mathsf{p}], \mathbb{P}, \mathbb{M}, \mathbb{G}, \{\mathsf{r}\}) = \mathsf{s}[\mathsf{p}]?\ell_1(\mathsf{x}) \mid \mathsf{s}[\mathsf{p}]!\ell_4(7)$ if $\mathbb{P} = \mathsf{s}[\mathsf{p}]?\ell_1(x) \mid \mathsf{s}[\mathsf{p}]?\ell_2(y) \mid \mathsf{s}[\mathsf{p}]!\ell_3(\mathsf{true})$, $\mathbb{M} = \mathsf{q}?\ell_1(\mathsf{Int}) \mid \mathsf{r}?\ell_2(\mathsf{Int}) \mid \mathsf{q}'!\ell_3(\mathsf{Bool})$, $\mathbb{G} = \mathsf{p} \to \mathsf{q}' : \ell_4(\mathsf{Int})$, and $\mathscr{P}$ contains $y!\ell_4(7)$.

We can now define the adaptation rule:

$$\frac{F(\sigma) = (\mathbb{G}, \mathscr{K}, \sigma') \quad \forall \mathsf{p} \in \mathsf{pa}(\mathbb{G}) \setminus \mathscr{A}.\mathbb{Q}_\mathsf{p} \in \mathscr{P} \;\&\; \mathbb{Q}_\mathsf{p} \propto \hat{\mathbb{G}} \upharpoonright \mathsf{p} \quad \mathbb{M}'_\mathsf{p} = \mathsf{mon}(\mathsf{p}, \mathbb{M}_\mathsf{p}, \hat{\mathbb{G}}, \mathscr{K}) \quad \mathbb{P}'_\mathsf{p} = \mathsf{proc}(\mathsf{s}[\mathsf{p}], \mathbb{P}_\mathsf{p}, \mathbb{M}_\mathsf{p}, \hat{\mathbb{G}}, \mathscr{K})}{(\nu\, \mathsf{s})\,(\prod_{\mathsf{p} \in \mathscr{A}} \mathbb{M}_\mathsf{p}[\mathbb{P}_\mathsf{p}]) \between \sigma \longrightarrow (\nu\, \mathsf{s})\,(\prod_{\mathsf{p} \in \mathscr{A} \setminus \mathscr{K}} \mathbb{M}'_\mathsf{p}[\mathbb{P}'_\mathsf{p}]) \parallel \prod_{\mathsf{p} \in \mathsf{pa}(\mathbb{G}) \setminus \mathscr{A}} \hat{\mathbb{G}} \upharpoonright \mathsf{p}\,[\mathbb{Q}_\mathsf{p}\{\mathsf{s}[\mathsf{p}]/y\}]) \between \sigma'}\ \textsc{Adapt}$$

Rule [ADAPT] must be used when the adaptation function F applied to the control data σ is defined. In this case F returns a global type $\mathbb{G}$, a set of participants $\mathscr{K}$ which must be killed and a new control data σ'. The global type $\hat{\mathbb{G}}$ is $\mathbb{G}$ with possibly some relabelling, as defined by rule [ADAPT]. The set of participants before the adaptation is $\mathscr{A}$. After the adaptation the set of participants is $(\mathscr{A} \setminus \mathscr{K}) \cup (\mathsf{pa}(\mathbb{G}) \setminus \mathscr{A})$, i.e. the initial set of participants without the killed ones, plus the participants of $\mathbb{G}$ which do not occur in $\mathscr{A}$. For these new participants the processes are taken from the collection $\mathscr{P}$ as in rule [INITS]. Instead for the participants in $\mathscr{A} \setminus \mathscr{K}$ we compute the new monitors and processes using the mappings mon and proc, respectively.

Table 5 lists the reduction rules of networks and systems, which are not discussed above. Evaluation contexts are defined by $\mathscr{E} ::= [\,] \mid \mathscr{E} \parallel N \mid (\nu s)\mathscr{E}$.

A system starts reducing session initiators. Then, in each reduction step, rule [ADAPT] is used whenever the adaptation function applied to the control data is defined. Otherwise the application of rule [OP] has precedence over the remaining rules for each monitored process.

A system is *initial* if it is the parallel compositions of news. Starting from an initial system the reduction rules produce only monitored processes $\mathbb{M}[\mathbb{P}]$ which satisfy $\mathbb{P} \propto \mathbb{M}$. The crucial case is rule [ADAPT], where the mapping proc builds the processes so that the adequacy between processes and monitors can be guaranteed.

Theorem 5.1 *If $\mathscr{S}$ is an initial system and $\mathscr{S} \longrightarrow \mathscr{S}'$, then all monitored processes in $\mathscr{S}'$ satisfy the adequacy condition.*

The proof of this theorem requires typing rules for networks and does not fit in this paper.

The closure $\mathscr{C}(\Sigma)$ of a set Σ of monitors is the smallest set which contains Σ and such that:

- $\mathsf{M} \in \mathscr{C}(\Sigma)$ implies $\mathsf{M} \setminus \mathscr{K} \in \mathscr{C}(\Sigma)$ for an arbitrary $\mathscr{K}$;
- $\mathbb{M} \in \mathscr{C}(\Sigma)$ and $\mathbb{M}' \in \mathscr{C}(\Sigma)$ imply $\mathbb{M} \mid \mathbb{M}' \in \mathscr{C}(\Sigma)$.

We can also guarantee progress of systems under a natural condition on the collection $\mathscr{P}$. A collection $\mathscr{P}$ is *complete for a given set of global types $\mathscr{G}$* if, for every $\mathbb{G} \in \mathscr{G}$, there are processes in $\mathscr{P}$ for the closure of the set of monitors obtained from $\mathbb{G}$ by projecting onto participants.

Theorem 5.2 *If $\mathscr{P}$ is complete for $\mathscr{G}$, the adaptation function F only produces global types in $\mathscr{G}$ and $\mathscr{S}$ is an initial system with global types in $\mathscr{G}$, then $\mathscr{S}$ has progress.*

The proof of this theorem uses previous theorem and the technique illustrated in [7].

$$\frac{\mathbb{P} \xrightarrow{\tau} \mathbb{P}'}{\mathbb{M}[\mathbb{P}] \longrightarrow \mathbb{M}[\mathbb{P}']}\ \textsc{Tau} \qquad \frac{N_1 \equiv N'_1 \quad N'_1 \longrightarrow N'_2 \quad N_2 \equiv N'_2}{N_1 \longrightarrow N_2}\ \textsc{Equiv}$$

$$\frac{N \longrightarrow N'}{\mathscr{E}[N] \between \sigma \longrightarrow \mathscr{E}[N'] \between \sigma}\ \textsc{SN} \qquad \frac{N \between \sigma \longrightarrow N' \between \sigma'}{\mathscr{E}[N] \between \sigma \longrightarrow \mathscr{E}[N'] \between \sigma'}\ \textsc{Ctx}$$

Table 5: Further network and system reductions.

6 Conclusion and Related Works

To the best of our knowledge the present paper provides the first formal model of self-adapting multiparty sessions in which control data trigger the system reconfigurations.

As already mentioned, our work builds on [8], where a calculus based on global types, monitors and processes similar to the present one was introduced. There are two main points of departure from that work. First, in the calculus of [8], the global type decides *when* the adaptation can take place, since its monitors prescribe when some participants have to check data and then send a request of adaptation to the other participants. Second, in [8] any adaptation involves all session participants, by requiring a global new reconfiguration of the whole system. In contrast, in our calculus adaptation is only triggered by dynamic changes in control data, so giving rise to unpredictable reconfiguration steps. Furthermore, reconfiguration can involve a partial set of communications only, when a part of the behaviour must be changed. The main technical novelty allowing these features is the parallel composition of monitors. Moreover, communication is synchronous here and asynchronous in [8]. We adopted this choice with the goal of designing a calculus as simple as possible, while capturing main issues concerning a safe self-reconfiguration of the participant behaviours in unuspected adaptations. Using asynchronous communications would add the technical complication of reconfiguring channel queues during adaptation, which is left for future work.

Works addressing adaptation for multiparty communications include [1], [4], [9] and [6]. In paper [1] global and session types are used to guarantee deadlock-freedom in a calculus of multiparty sessions with asynchronous communications. Only part of the running code is updated. Two different conditions are given for ensuring liveness. The first condition requires that all channel queues are empty before updating. The second condition requires a partial order between the session participants with a unique minimal element. The participants are then updated following this order. We do not need such conditions for progress, since communication is synchronous and adaptation is done in one single big step. The work [4] enhances a choreographic language with constructs defining adaptation scopes and dynamic code update; an associated endpoint language for local descriptions, and a projection mechanism for obtaining (low-level) endpoint specifications from (high-level) choreographies are also described. A typing discipline for these languages is left for future work. The paper [9] proposes a choreographic language for distributed applications. Adaptation follows a rule-based approach, in which all interactions, under all possible changes produced by the adaptation rules, proceed as prescribed by an abstract model. In particular, the system is deadlock-free by construction. The adaptive system is composed by interacting participants deployed on different locations, each executing its own code. The calculus of [6] is inspired by [8], the main difference being that security violations trigger adaptation mechanisms that prevent violations to occur and/or to propagate their effect in the systems.

There is substantial difference between our monitors and run-time enforcement monitors of [15], where monitors terminate the execution if it is about to violate the security policy being enforced.

Future work includes the addition of a local state to each participant. This can allow us to model internal interactions, based on local data, among single threaded processes whose parallel composition implements a participant. Using both internal interactions and communications among participants enables the application of our calculus to more realistic cases of study. In general, we plan to consider application-driven instances of our global adaptation mechanisms, in which concrete adaptation functions are used.

Finally, we intend to investigate whether our approach can be useful in providing a formal model to mechanisms of code hot-swapping in programming languages (such as Erlang).

Acknowledgments We are grateful to the anonymous reviewers for their useful suggestions and remarks.

References

[1] Gabrielle Anderson & Julian Rathke (2012): *Dynamic Software Update for Message Passing Programs*. In: *APLAS'12*, *LNCS* 7705, Springer, pp. 207–222, doi:10.1007/978-3-642-35182-2_15.

[2] Lorenzo Bettini, Mario Coppo, Loris D'Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini & Nobuko Yoshida (2008): *Global Progress in Dynamically Interleaved Multiparty Sessions*. In: *CONCUR'08*, *LNCS* 5201, Springer, pp. 418–433, doi:10.1007/978-3-540-85361-9_33.

[3] Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda & Nobuko Yoshida (2013): *Monitoring Networks through Multiparty Session Types*. In: *FMOODS/FORTE'13*, *LNCS* 7892, Springer, pp. 50–65, doi:10.1007/978-3-642-38592-6_5.

[4] Mario Bravetti, Marco Carbone, Thomas T. Hildebrandt, Ivan Lanese, Jacopo Mauro, Jorge A. Pérez & Gianluigi Zavattaro (2014): *Towards Global and Local Types for Adaptation*. In: *SEFM'13*, *LNCS* 8368, Springer, pp. 3–14, doi:10.1007/978-3-319-05032-4_1.

[5] Marco Carbone, Kohei Honda & Nobuko Yoshida (2012): *Structured Communication-Centered Programming for Web Services*. *ACM Transactions on Programming Languages and Systems* 34(2), pp. 8:1–8:78, doi:10.1145/2220365.2220367.

[6] Ilaria Castellani, Mariangiola Dezani-Ciancaglini & Jorge A. Pérez (2014): *Self-Adaptation and Secure Information Flow in Multiparty Structured Communications: A Unified Perspective*. In: *BEAT'14*, *EPTCS* 162, Open Publishing Association, pp. 9–18, doi:10.4204/EPTCS.162.2.

[7] Mario Coppo, Mariangiola Dezani-Ciancaglini, Luca Padovani & Nobuko Yoshida (2015): *A Gentle Introduction to Multiparty Asynchronous Session Types*. In: *SFM'15*, *LNCS* 9104, Springer, pp. 146–178, doi:10.1007/978-3-319-18941-3_4.

[8] Mario Coppo, Mariangiola Dezani-Ciancaglini & Betti Venneri (2015): *Self-Adaptive Multiparty Sessions*. *Service Oriented Computing and Applications* 9(3-4), pp. 249–268, doi:10.1007/s11761-014-0171-9.

[9] Mila Dalla Preda, Saverio Giallorenzo, Ivan Lanese, Jacopo Mauro & Maurizio Gabbrielli (2014): *AIOCJ: A Choreographic Framework for Safe Adaptive Distributed Applications*. In: *SLE'14*, *LNCS* 8706, Springer, pp. 161–170, doi:10.1007/978-3-319-11245-9_9.

[10] Pierre-Malo Deniélou & Nobuko Yoshida (2011): *Dynamic Multirole Session Types*. In: *POPL'11*, ACM Press, pp. 435–446, doi:10.1145/1926385.1926435.

[11] Kohei Honda, Vasco T. Vasconcelos & Makoto Kubo (1998): *Language Primitives and Type Disciplines for Structured Communication-based Programming*. In: *ESOP'98*, *LNCS* 1381, Springer, pp. 22–138, doi:10.1007/BFb0053567.

[12] Kohei Honda, Nobuko Yoshida & Marco Carbone (2008): *Multiparty Asynchronous Session Types*. In: *POPL'08*, ACM Press, pp. 273–284, doi:10.1145/1328438.1328472.

[13] Hans Hüttel, Ivan Lanese, Vasco Thudichum Vasconcelos, Luís Caires, Marco Carbone, Pierre-Malo Deniélou, Dimitris Mostrous, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira & Gianluigi Zavattaro (2014): *Foundations of Behavioural Types*. http://www.behavioural-types.eu/publications/WG1-State-of-the-Art.pdf.

[14] Benjamin C. Pierce (2002): *Types and Programming Languages*. MIT Press.

[15] Fred B. Schneider (2000): *Enforceable Security Policies*. *ACM Transactions on Information and System Security* 3(1), pp. 30–50, doi:10.1145/353323.353382.